

On the half-Cauchy prior for a global scale parameter

Nicholas G. Polson*
 University of Chicago

James G. Scott†
 The University of Texas at Austin

First draft: June 2010
 This draft: September 2011

Abstract

This paper argues that the half-Cauchy distribution should replace the inverse-Gamma distribution as a default prior for a top-level scale parameter in Bayesian hierarchical models, at least for cases where a proper prior is necessary. Our arguments involve a blend of Bayesian and frequentist reasoning, and are intended to complement the original case made by Gelman (2006) in support of the folded- t family of priors. First, we generalize the half-Cauchy prior to the wider class of hypergeometric inverted-beta priors. We derive expressions for posterior moments and marginal densities when these priors are used for a top-level normal variance in a Bayesian hierarchical model. We go on to prove a proposition that, together with the results for moments and marginals, allows us to characterize the frequentist risk of the Bayes estimators under all global-shrinkage priors in the class. These theoretical results, in turn, allow us to study the frequentist properties of the half-Cauchy prior versus a wide class of alternatives. The half-Cauchy occupies a sensible “middle ground” within this class: it performs very well near the origin, but does not lead to drastic compromises in other parts of the parameter space. This provides an alternative, classical justification for the repeated, routine use of this prior. We also consider situations where the underlying mean vector is sparse, where we argue that the usual conjugate choice of an inverse-gamma prior is particularly inappropriate, and can lead to highly distorted posterior inferences. Finally, we briefly summarize some open issues in the specification of default priors for scale terms in hierarchical models.

Keywords: hierarchical models; normal scale mixtures; shrinkage

Acknowledgements: the authors would like to thank an anonymous referee, the editor, and the associate editor, whose many useful suggestions have helped us to improve the manuscript.

*nicholas.polson@chicagobooth.edu

†james.scott@mccombs.utexas.edu

1 Introduction

Consider a normal hierarchical model where, for $i = 1, \dots, p$,

$$\begin{aligned} (y_i | \beta_i, \sigma^2) &\sim N(\beta_i, \sigma^2) \\ (\beta_i | \lambda^2, \sigma^2) &\sim N(0, \lambda^2 \sigma^2) \\ \lambda^2 &\sim g(\lambda^2). \end{aligned}$$

This prototype case embodies a very general problem in Bayesian inference: how to choose default priors for top-level variances (here λ^2 and σ^2) in a hierarchical model.

The routine use of Jeffreys' prior for the error variance, $p(\sigma^2) \propto \sigma^{-2}$, poses no practical issues. This is not the case for $p(\lambda^2)$, however, as the improper prior $p(\lambda^2) \propto \lambda^{-2}$ leads to an improper posterior. This can be seen from the marginal likelihood:

$$p(y | \lambda^2) \propto \prod_{i=1}^p (1 + \lambda^2)^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \sum_{i=1}^p \frac{y_i^2}{1 + \lambda^2}\right),$$

where we have taken $\sigma^2 = 1$ for convenience. This is positive at $\lambda^2 = 0$; therefore, whenever the prior $p(\lambda^2)$ fails to be integrable at the origin, so too will the posterior. A number of default choices have been proposed to overcome this issue. A classic reference is Tiao and Tan (1965); a very recent one is Morris and Tang (2011), who use a flat prior $p(\lambda^2) \propto 1$.

We focus on a proposal by Gelman (2006), who studies the class of half- t priors for the scale parameter λ :

$$p(\lambda | d) \propto \left(1 + \frac{\lambda^2}{d}\right)^{-(d+1)/2}$$

for some degrees-of-freedom parameter d . The half- t prior has the appealing property that its density evaluates to a nonzero constant at $\lambda = 0$. This distinguishes it from the usual conjugate choice of an inverse-gamma prior for λ^2 , whose density vanishes at $\lambda = 0$. As Gelman (2006) points out, posterior inference under these priors is no more difficult than it is under an inverse-gamma prior, using the simple trick of parameter expansion.

These facts lead to a simple, compelling argument against the use of the inverse-gamma prior for variance terms in models such as that above. Since the marginal likelihood of the data, considered as a function of λ , does not vanish when $\lambda = 0$, neither should the prior density $p(\lambda)$. Otherwise, the posterior distribution for λ will be inappropriately biased away from zero. This bias, moreover, is most severe near the origin, precisely in the region of parameter space where the benefits of shrinkage become most pronounced.

This paper studies the special case of a half-Cauchy prior for λ with three goals in mind. First, we embed it in the wider class of hypergeometric inverted-beta priors for λ^2 , and derive expressions for the resulting posterior moments and marginal densities. Second, we derive expressions for the classical risk of Bayes estimators arising from this class of priors. In particular, we prove a result that allows us to characterize the improvements in risk near the origin ($\|\beta\| \approx 0$) that are possible using the wider class. Having proven our risk results for all members of this wider class, we then return to the special case of the half-Cauchy; we find that the frequentist risk profile of the resulting Bayes estimator is quite

favorable, and rather similar to that of the positive-part James–Stein estimator. Therefore Bayesians can be comfortable using the prior on purely frequentist grounds.

Third, we attempt to provide some insight about the use of such priors in situations where β is expected to be sparse. We find that the arguments of Gelman (2006) in favor of the half-Cauchy are, if anything, amplified in the presence of sparsity, and that the inverse-gamma prior can have an especially distorting effect on posterior inference for sparse signals.

Overall, our results provide a complementary set of arguments in addition to those of Gelman (2006) that support the routine use of the half-Cauchy prior: its excellent (frequentist) risk properties, and its sensible behavior in the presence of sparsity compared to the usual conjugate alternative. Bringing all these arguments together, we contend that the half-Cauchy prior is a sensible default choice for a top-level variance in Gaussian hierarchical models. We echo the call for it to replace inverse-gamma priors in routine use, particularly given the availability of a simple parameter-expanded Gibbs sampler for posterior computation.

2 Inverted-beta priors and their generalizations

Consider the family of inverted-beta priors for λ^2 :

$$p(\lambda^2) = \frac{(\lambda^2)^{b-1} (1+\lambda^2)^{-(a+b)}}{\text{Be}(a,b)},$$

where $\text{Be}(a,b)$ denotes the beta function, and where a and b are positive reals. A half-Cauchy prior for λ corresponds to an inverted-beta prior for λ^2 with $a = b = 1/2$. This family also generalizes the robust priors of Strawderman (1971) and Berger (1980); the normal-exponential-gamma prior of Griffin and Brown (2005); and the horseshoe prior of Carvalho et al. (2010). The inverted-beta distribution is also known as the beta-prime or Pearson Type VI distribution. An inverted-beta random variable is equal in distribution to the ratio of two gamma-distributed random variables having shape parameters a and b , respectively, along with a common scale parameter.

The inverted-beta family is itself a special case of a new, wider class of hypergeometric inverted-beta distributions having the following probability density function:

$$p(\lambda^2) = C^{-1} (\lambda^2)^{b-1} (\lambda^2 + 1)^{-(a+b)} \exp\left\{-\frac{s}{1+\lambda^2}\right\} \left\{\tau^2 + \frac{1-\tau^2}{1+\lambda^2}\right\}^{-1}, \quad (1)$$

for $a > 0$, $b > 0$, $\tau^2 > 0$, and $s \in \mathcal{R}$. This comprises a wide class of priors leading to posterior moments and marginals that can be expressed using confluent hypergeometric functions. In Appendix A we give details of these computations, which yield

$$C = e^{-s} \text{Be}(a,b) \Phi_1(b, 1, a+b, s, 1 - 1/\tau^2), \quad (2)$$

where Φ_1 is the degenerate hypergeometric function of two variables (Gradshteyn and Ryzhik, 1965, 9.261). This function can be calculated accurately and rapidly by transforming it into a convergent series of ${}_2F_1$ functions (§9.2 of Gradshteyn and Ryzhik, 1965; Gordy, 1998), making evaluation of (2) quite fast for most choices of the parameters.

Both τ and s are global scale parameters, and do not control the behavior of $p(\lambda)$ at 0 or ∞ . The parameters a and b are analogous to those of beta distribution. Smaller values of a encourage heavier tails in $p(\beta)$, with $a = 1/2$, for example, yielding Cauchy-like tails. Smaller values of b encourage $p(\beta)$ to have more mass near the origin, and eventually to become unbounded; $b = 1/2$ yields, for example, $p(\beta) \approx \log(1 + 1/\beta^2)$ near 0.

We now derive expressions for the moments of $p(\beta | \mathbf{y}, \sigma^2)$ and the marginal likelihood $p(\mathbf{y} | \sigma^2)$ for priors in this family. As a special case, we easily obtain the posterior mean for β under a half-Cauchy prior on λ .

Given λ^2 and σ^2 , the posterior distribution of β is multivariate normal, with mean m and variance V given by

$$m = \left(1 - \frac{1}{1 + \lambda^2}\right) \mathbf{y} \quad , \quad V = \left(1 - \frac{1}{1 + \lambda^2}\right) \sigma^2.$$

Define $\kappa = 1/(1 + \lambda^2)$. By Fubini's theorem, the posterior mean and variance of β are

$$\mathbb{E}(\beta | \mathbf{y}, \sigma^2) = \{1 - \mathbb{E}(\kappa | \mathbf{y}, \sigma^2)\} \mathbf{y} \quad (3)$$

$$\text{var}(\beta | \mathbf{y}, \sigma^2) = \{1 - \mathbb{E}(\kappa | \mathbf{y}, \sigma^2)\} \sigma^2, \quad (4)$$

now conditioning only on σ^2 .

It is most convenient to work with $p(\kappa)$ instead:

$$p(\kappa) \propto \kappa^{a-1} (1-\kappa)^{b-1} \left\{ \frac{1}{\tau^2} + \left(1 - \frac{1}{\tau^2}\right) \kappa \right\}^{-1} e^{-\kappa s}. \quad (5)$$

The joint density for κ and \mathbf{y} takes the same functional form:

$$p(y_1, \dots, y_p, \kappa) \propto \kappa^{a'-1} (1-\kappa)^{b-1} \left\{ \frac{1}{\tau^2} + \left(1 - \frac{1}{\tau^2}\right) \kappa \right\}^{-1} e^{-\kappa s'},$$

with $a' = a + p/2$, and $s' = s + Z/2\sigma^2$ for $Z = \sum_{i=1}^p y_i^2$. Hence the posterior for λ^2 is also a hypergeometric inverted-beta distribution, with parameters (a', b, τ^2, s') .

Next, the moment-generating function of (5) is easily shown to be

$$M(t) = e^t \frac{\Phi_1(b, 1, a+b, s-t, 1-1/\tau^2)}{\Phi_1(b, 1, a+b, s, 1-1/\tau^2)}.$$

See, for example, Gordy (1998). Expanding Φ_1 as a sum of ${}_1F_1$ functions and using the differentiation rules given in Chapter 15 of Abramowitz and Stegun (1964) yields

$$\mathbb{E}(\kappa^n | \mathbf{y}, \sigma^2) = \frac{(a')_n}{(a'+b)_n} \frac{\Phi_1(b, 1, a'+b+n, s', 1-1/\tau^2)}{\Phi_1(b, 1, a'+b, s', 1-1/\tau^2)}. \quad (6)$$

Combining the above expression with (3) and (4) yields the conditional posterior mean and variance for β , given \mathbf{y} and σ^2 . Similarly, the marginal density $p(\mathbf{y} | \sigma^2)$ is a simple expression involving the ratio of prior to posterior normalizing constants:

$$p(\mathbf{y} | \sigma^2) = (2\pi\sigma^2)^{-p/2} \exp\left(-\frac{Z}{2\sigma^2}\right) \frac{\text{Be}(a', b)}{\text{Be}(a, b)} \frac{\Phi_1(b, 1, a'+b, s', 1-1/\tau^2)}{\Phi_1(b, 1, a+b, s, 1-1/\tau^2)}.$$

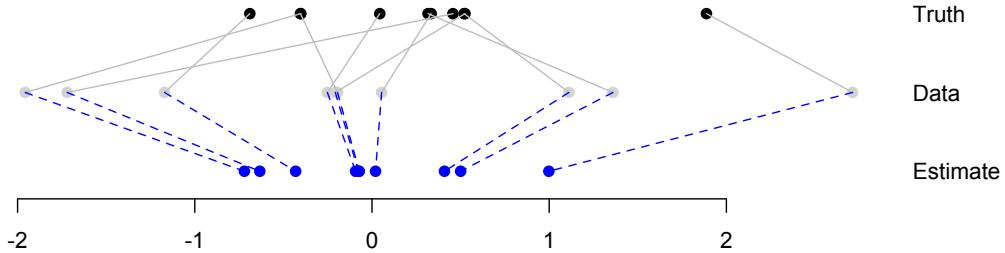


Figure 1: Ten true means drawn from a standard normal distribution; data from these means under standard normal noise; shrinkage estimates under the half-Cauchy prior for λ .

3 Classical risk results

These priors are useful in situations where standard priors like the inverse-gamma or Jeffreys' are inappropriate or ill-behaved. Non-Bayesians will find them useful for generating easily computable shrinkage estimators that have known risk properties. Bayesians will find them useful for generating computationally tractable priors for a variance parameter. We argue that these complementary but overlapping goals can both be satisfied for the special case of the half-Cauchy. To show this, we first characterize the risk properties of the Bayes estimators that result from the wider family of priors used for a normal mean under a quadratic loss. Our analysis shows that:

1. The hypergeometric–beta family provides a large class of Bayes estimators that will perform no worse than the MLE in the tails, i.e. when $\|\beta\|^2$ is large.
2. Major improvements over the James–Stein estimator are possible near the origin. This can be done in several ways: by choosing a large relative to b , by choosing a and b both less than 1, by choosing s negative, or by choosing $\tau < 1$. Each of these choices involves a compromise somewhere else in the parameter space.

We now derive expressions for the classical risk, as a function of $\|\beta\|$, for the resulting Bayes estimators under hypergeometric inverted-beta priors. Assume without loss of generality that $\sigma^2 = 1$, and let $p(\mathbf{y}) = \int p(\mathbf{y}|\beta)p(\beta)d\beta$ denote the marginal density of the data. Following Stein (1981), write the mean-squared error of the posterior mean $\hat{\beta}$ as

$$E(\|\hat{\beta} - \beta\|^2) = p + E_{\mathbf{y}} \left(\|g(\mathbf{y})\|^2 + 2 \sum_{i=1}^p \frac{\partial}{\partial y_i} g(\mathbf{y}) \right),$$

where $g(\mathbf{y}) = \nabla \log p(\mathbf{y})$. In turn this can be written as

$$E(\|\hat{\beta} - \beta\|^2) = p + 4E_{\mathbf{y}|\beta} \left(\frac{\nabla^2 \sqrt{p(\mathbf{y})}}{\sqrt{p(\mathbf{y})}} \right).$$

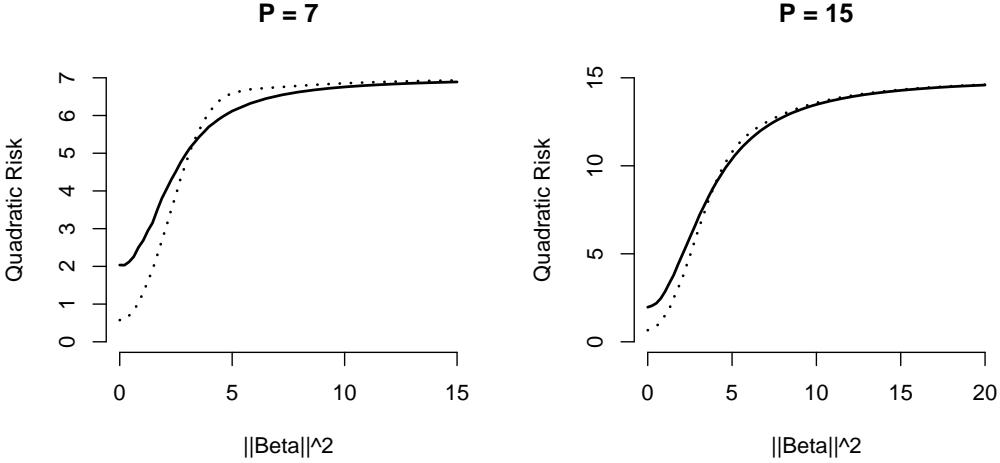


Figure 2: Mean-squared error as a function of $\|\beta\|$ for $p = 7$ and $p = 15$. Solid line: James–Stein estimator. Dotted line: Bayes estimator under a half-Cauchy prior for λ .

We now state our main result concerning computation of this quantity.

Proposition 1. Suppose that $\beta \sim N_p(0, \lambda^2 I)$, that $\kappa = 1/(1 + \lambda^2)$, and that the prior $p(\kappa)$ is such that $\lim_{\kappa \rightarrow 0,1} \kappa(1 - \kappa)p(\kappa) = 0$. Define

$$m_p(Z) = \int_0^1 \kappa^{\frac{p}{2}} e^{-\frac{Z}{2}\kappa} p(\kappa) d\kappa$$

for $Z = \sum_{i=1}^p y_i^2$. Then as a function of β , the quadratic risk of the posterior mean under $p(\kappa)$ is

$$E(\|\hat{\beta} - \beta\|^2) = p + 2E_{Z|\beta} \left\{ Z \frac{m_{p+4}(Z)}{m_p(Z)} - pg(Z) - \frac{Z}{2}g(Z)^2 \right\}, \quad (7)$$

where $g(Z) = E(\kappa | Z)$, and where

$$Z \frac{m_{p+4}(Z)}{m_p(Z)} = (p + Z + 4)g(Z) - (p + 2) - E_{\kappa|Z} \left\{ 2\kappa(1 - \kappa) \frac{p'(\kappa)}{p(\kappa)} \right\}. \quad (8)$$

Proof. See Appendix B. □

Proposition 1 is useful because it characterizes the risk in terms of two known quantities: the integral $m_p(Z)$, and the posterior expectation $g(Z) = E(\kappa | Z)$. Using the results of the previous section, these are easily obtained under a hypergeometric inverted-beta prior for λ^2 . Furthermore, given $\|\beta\|$, $Z = U^2 + V$ in distribution, where

$$\begin{aligned} U &\sim N(\|\beta\|, 1) \\ V &\sim \chi_{p-1}^2. \end{aligned}$$

The risk of the Bayes estimator is therefore easy to evaluate as a function of $\|\beta\|^2$. These expressions can be compared to those of, for example, George et al. (2006), who consider Kullback–Leibler predictive risk for similar priors.

Our interest is in the special case $a = b = 1/2$, $\tau = 1$, and $s = 0$, corresponding to a half-Cauchy prior for the global scale λ . Figure 2 shows the classical risk of the Bayes estimator under this prior for $p = 7$ and $p = 15$. The risk of the James–Stein estimator is shown for comparison. These pictures look similar for other values of p , and show overall that the half-Cauchy prior for λ leads to a Bayes estimator that is competitive with the James–Stein estimator.

Figure 1 shows a simple example of the posterior mean under the half-Cauchy prior for λ when $p = 10$, calculated for fixed σ using the results of the previous section. For this particular value of β the expected squared-error risk of the MLE is 10, and the expected squared-error risk of the half-Cauchy posterior mean is 8.6.

A natural question is: of all the hypergeometric inverted-beta priors, why choose the half-Cauchy? There is no iron-clad reason to do so, of course, and we can imagine many situations where subjective information would support a different choice. But in examining many other members of the class, we have observed that the half-Cauchy seems to occupy a sensible “middle ground” in terms of frequentist risk. To study this, we are able to appeal to the theory of the previous section. See, for example, Figure 3, which compares several members of the class for the case $p = 7$. Observe that large gains over James–Stein near the origin are possible, but only at the expense of minimaxity. The half-Cauchy, meanwhile, still improves upon the James–Stein estimator near the origin, but does not sacrifice good risk performance in other parts of the parameter space. From a purely classical perspective, it looks like a sensible default choice, suitable for repeated general use.

4 Global scale parameters in local-shrinkage models

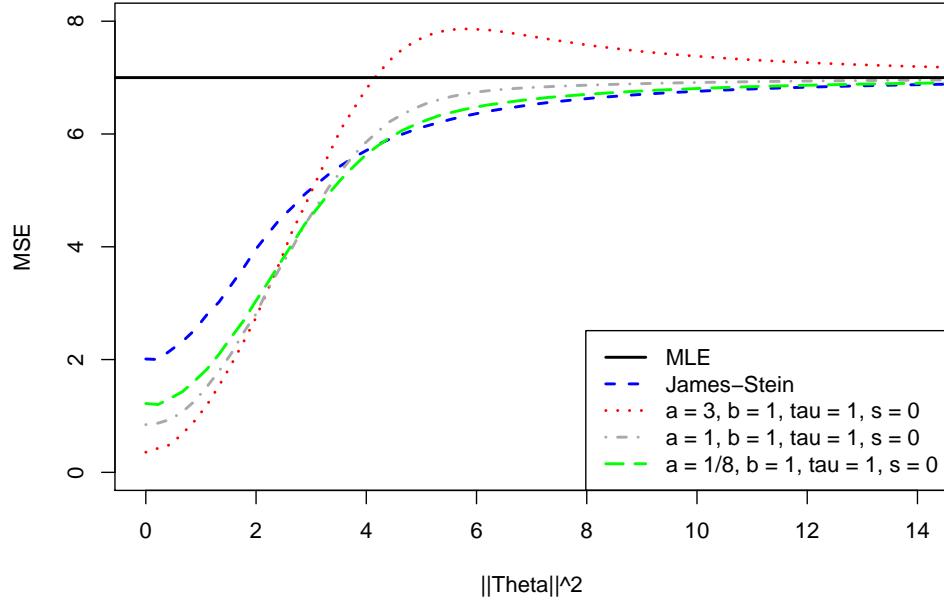
A now-canonical modification of the basic hierarchical model from the introduction involves the use of local shrinkage parameters:

$$\begin{aligned} (y_i | \beta_i, \sigma^2) &\sim N(\beta_i, \sigma^2) \\ (\beta_i | \lambda^2, u_i^2, \sigma^2) &\sim N(0, \lambda^2 \sigma^2 u_i^2) \\ u_i^2 &\sim f(u_i^2) \\ \lambda^2 &\sim g(\lambda^2). \end{aligned}$$

Mixing over u_i leads to a non-Gaussian marginal for β_i . For example, choosing an exponential prior for each u_i^2 results in a Laplace (lasso-type) prior. This class of models provides a Bayesian alternative to penalized-likelihood estimation. When the underlying vector of means is sparse, these global-local shrinkage models can lead to large improvements in both estimation and prediction compared with pure global shrinkage rules. There is a large literature on the choice of $p(u_i^2)$, with Polson and Scott (2011a) providing a recent review.

As many authors have documented, strong global shrinkage combined with heavy-tailed local shrinkage is why these sparse Bayes estimators work so well at sifting signals from

Changing a , $p = 7$



Changing s , $p = 7$

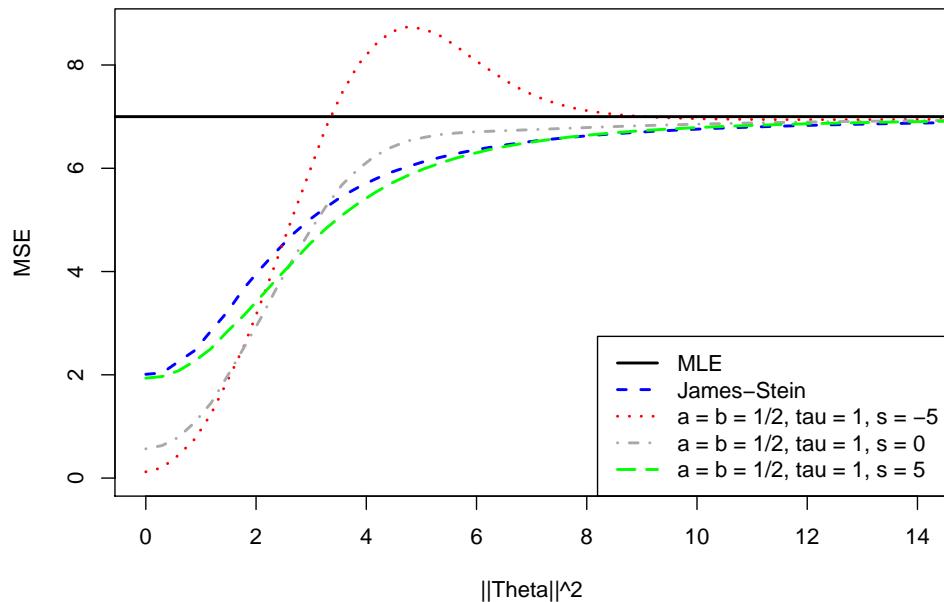


Figure 3: Mean-squared error as a function of $\|\beta\|^2$ for $p = 7$ and various cases of the hypergeometric inverted-beta hyperparameters.

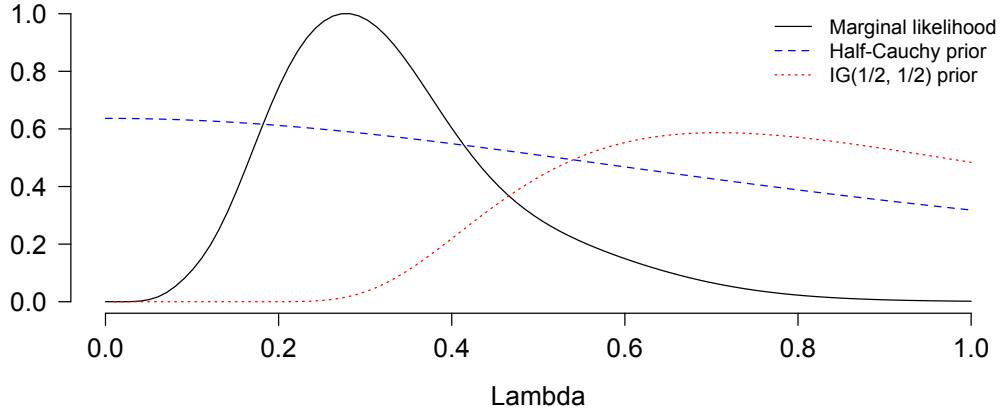


Figure 4: The black line shows the marginal likelihood of the data as a function of λ under a horseshoe prior for each β_i . (The likelihood has been renormalized to obtain a maximum of 1.) The blue and red lines show two priors for λ : the half-Cauchy, and that induced by an inverse-Gamma prior on λ^2 .

noise. Intuitively, the idea is that λ acts as a global parameter that adapts to the underlying sparsity of the signal. When few signals are present, it is quite common for the marginal likelihood of \mathbf{y} as a function of λ to concentrate near 0 (“shrink globally”), and for the signals to be flagged via very large values of the local shrinkage parameters u_i^2 (“act locally”). Indeed, in some cases the marginal maximum-likelihood solution can be the degenerate $\hat{\lambda} = 0$ (see, for example, Tiao and Tan, 1965).

The classical risk results of the previous section no longer apply to a model with these extra local-shrinkage parameters, since the marginal distribution of β , given λ , is not multivariate normal. Nonetheless, the case of sparsity serves only to amplify the purely Bayesian argument in favor of the half-Cauchy prior for a global scale parameter—namely, the argument that $p(\lambda | \mathbf{y})$ should not be artificially pulled away from zero by an inverse-gamma prior.

Figure 4 vividly demonstrates this point. We simulated data from a sparse model where β contained the entries $(5, 4, 3, 2, 1)$ along with 45 zeroes, and where $y_{ij} \sim N(0, 1)$ for $j = 1, 2, 3$. We then used Markov Chain Monte Carlo to compute the marginal likelihood of the data as a function of λ , assuming that each β_i has a horseshoe prior (Carvalho et al., 2010). This can be approximated by assuming a flat prior for λ (here truncated above at 10), and then computing the conditional likelihood $p(\mathbf{y} | \lambda, \sigma, u_1^2, \dots, u_p^2)$ over a discrete grid of λ values at each step of the Markov chain. The marginal likelihood function can then be approximated as the pointwise average of the conditional likelihood over the samples from the joint posterior.

This marginal likelihood has been renormalized to obtain a maximum of 1 and then plotted alongside two alternatives: a half-Cauchy prior for λ , and the prior induced by assuming that $\lambda^2 \sim \text{IG}(1/2, 1/2)$. Under the inverse-gamma prior, there will clearly be an

inappropriate biasing of $p(\lambda)$ away from zero, which will negatively affect the ability of the model to handle sparsity efficiently. For data sets with even more “noise” entries in \mathbf{y} , the distorting effect of a supposedly “default” inverse-gamma prior will be even more pronounced, as the marginal likelihood will favor values of τ very near zero (along with a small handful of very large u_i^2 terms).

5 Discussion

On strictly Bayesian grounds, the half-Cauchy is a sensible default prior for scale parameters in hierarchical models: it tends to a constant at $\lambda = 0$; it is quite heavy-tailed; and it leads to simple conjugate MCMC routines, even in more complex settings. All these desirable features are summarized by Gelman (2006). Our results give a quite different, classical justification for this prior in high-dimensional settings: its excellent quadratic risk properties. The fact that two independent lines of reasoning both lead to the same prior is a strong argument in its favor as a default proper prior for a shared variance component. We also recommend scaling the β_i ’s by σ , as reflected in the hierarchical model from the introduction. This is the approach taken by Jeffreys (1961, Section 5.2), and we cannot improve upon his arguments.

In addition, our hypergeometric inverted-beta class provides a useful generalization of the half-Cauchy prior, in that it allows for greater control over global shrinkage through τ and s . It leads to a large family of estimators with a wide range of possible behavior, and generalizes the form noted by Maruyama (1999), which contains the positive-part James–Stein estimator as a limiting, improper case. Further study of this class may yield interesting frequentist results, quite apart from the Bayesian implications considered here. The expressions for marginal likelihoods also have connections with recent work on generalized g -priors (Maruyama and George, 2010; Polson and Scott, 2011b). Finally, all estimators arise from proper priors on λ^2 , and will therefore be admissible.

There are still many open issues in default Bayes analysis for hierarchical models that are not addressed by our results. One issue is whether to mix further over the scale in the half-Cauchy prior, $\lambda \sim C^+(0, \tau)$. One possibility here is simply to let $\tau \sim C^+(0, 1)$. We then get the following default “double” half-Cauchy prior for λ :

$$p(\lambda) = \frac{2}{\pi^2} \int_0^\infty \frac{1}{1 + \tau^2} \frac{1}{\tau(1 + \frac{\lambda^2}{\tau^2})} d\tau = \frac{\ln|\lambda|}{\lambda^2 - 1}.$$

Admittedly, it is difficult to know where to stop in this “turtles all the way down” approach to mixing over hyperparameters. (Why not, for example, mix still further over a scale parameter for τ ?) Even so, this prior has a number of appealing properties. It is proper, and therefore leads to a proper posterior; it is similar in overall shape to Jeffreys’ prior; and it is unbounded at the origin, and will therefore not down-weight the marginal likelihood as much as the half-Cauchy for near-sparse configurations of β . The implied prior on the shrinkage weight κ for the double half-Cauchy is

$$p(\kappa) \propto \frac{\ln(\frac{1-\kappa}{\kappa})}{1-2\kappa} \frac{1}{\sqrt{\kappa(1-\kappa)}}.$$

This is like the horseshoe prior on the shrinkage weight (Carvalho et al., 2010), but with an extra factor that comes from the fact that one is letting the scale itself be random with a $C^+(0, 1)$ prior.

We can also transform to the real line by letting $\psi = \ln \lambda^2$. For the half-Cauchy prior $p(\lambda) \propto 1/(1 + \lambda^2)$ this transformation leads to

$$p(\psi) \propto \frac{e^{\frac{\psi}{2}}}{1 + e^\psi} = \left(e^{\frac{\psi}{2}} + e^{-\frac{\psi}{2}} \right)^{-1} = \operatorname{sech}\left(\frac{\psi}{2}\right).$$

This is the hyperbolic secant distribution, which may provide fertile ground for further generalizations or arguments involving sensible choices for a default prior.

A more difficult issue concerns the prior scaling for λ in the presence of unbalanced designs—that is, when $y_{ij} \sim N(\beta_i, \sigma^2)$ for $j = 1, \dots, n_i$, and the n_i 's are not necessarily equal. In this case most formal non-informative priors for λ (e.g. the reference prior for a particular parameter ordering) involve complicated functions of the n_i 's (see, e.g. Yang and Berger, 1997). These expressions emerge from a particular mathematical formalism that, in turn, embodies a particular operational definition of “non-informative.”

We have focused on default priors that occupy a middle ground between formal non-informative analysis and pure subjective Bayes. This is clearly an important situation for the many practicing Bayesians who do not wish to use noninformative priors, whether for practical, mathematical, or philosophical reasons. An example of a situation in which formal noninformative priors for λ should not be used on mathematical grounds is when β is expected to be sparse; see Scott and Berger (2006) for a discussion of this issue in the context of multiple-testing. It is by no means obvious how, or even whether, the n_i 's should play a role in scaling λ within this (admittedly ill-defined) paradigm of “default” Bayes.

Finally, another open issue is the specification of default priors for scale parameters in non-Gaussian models. For example, in logistic regression, the likelihood is highly sensitive to large values of the underlying linear predictor. It is therefore not clear whether something so heavy-tailed as the half-Cauchy is an appropriate prior for the global scale term for logistic regression coefficients. All of these issues merit further research.

A Details for computing moments and marginals

The normalizing constant in (2) is

$$C = \int_0^1 \kappa^{\alpha-1} (1-\kappa)^{\beta-1} \left\{ \frac{1}{\tau^2} + \left(1 - \frac{1}{\tau^2} \right) \kappa \right\}^{-1} \exp(-s\kappa) d\kappa. \quad (9)$$

Let $\eta = 1 - \kappa$. Using the identity that $e^x = \sum_{m=0}^{\infty} x^m / m!$, we obtain

$$C = e^{-s} \sum_{m=0}^{\infty} \left[\frac{s^m}{m!} \int_0^1 \eta^{\beta+m-1} (1-\eta)^{\alpha-1} \{1 - (1 - 1/\tau^2)\eta\}^{-1} d\eta \right].$$

Using properties of the hypergeometric function ${}_2F_1$ (Abramowitz and Stegun, 1964, §15.1.1 and §15.3.1), this becomes, after some straightforward algebra,

$$C = e^{-s} \operatorname{Be}(\alpha, \beta) \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \frac{(\beta)_{m+n}}{(\alpha + \beta)_{m+n} m! n!} s^m (1 - 1/\tau^2)^m, \quad (10)$$

where $(a)_n$ is the rising factorial. Appendix C of Gordy (1998) proves that, for all $\alpha > 0$, $\beta > 0$, and $1/\tau^2 > 0$, the nested series in (10) converges to a positive real number, yielding

$$C = e^{-s} \operatorname{Be}(\alpha, \beta) \Phi_1(\beta, 1, \alpha + \beta, s, 1 - 1/\tau^2), \quad (11)$$

where Φ_1 is the degenerate hypergeometric function of two variables (Gradshteyn and Ryzhik, 1965, 9.261).

The Φ_1 function can be written as a double hypergeometric series,

$$\Phi_1(\alpha, \beta; \gamma; x, y) = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \frac{(\alpha)_{m+n} (\beta)_n}{(\gamma)_{m+n} m! n!} y^n x^m, \quad (12)$$

where $(c)_n$ is the rising factorial. We use three different representations of $\Phi_1(\alpha, \beta, \gamma, x, y)$ for handling different combinations of arguments, all from Gordy (1998). When $0 \leq y < 1$ and $x \geq 0$,

$$\Phi_1(\alpha, \beta, \gamma, x, y) = \sum_{n=0}^{\infty} \frac{(\alpha)_n x^n}{(\gamma)_n n!} {}_2F_1(\beta, \alpha + n; \gamma + n; y). \quad (13)$$

When $0 \leq y < 1$ and $x < 0$,

$$\Phi_1(\alpha, \beta, \gamma, x, y) = e^x \sum_{n=0}^{\infty} \frac{(\gamma - \alpha)_n (-x)^n}{(\gamma)_n n!} {}_2F_1(\beta, \alpha; \gamma + n; y). \quad (14)$$

Finally, when $y < 0$,

$$\Phi_1(\alpha, \beta, \gamma, x, y) = e^x (1 - y)^{-\beta} \Phi_1(\tilde{\alpha}, \beta, \gamma, -x, \tilde{y}), \quad (15)$$

where $\tilde{\alpha} = \gamma - \alpha$ and $\tilde{y} = y/(y - 1)$. Then either (13) or (14) may be used to evaluate the righthand side of (15), depending on the sign of x .

B Proof of Proposition 1

Proof. Begin with Stein's decomposition of risk. Following Equation (10) of Fourdrinier et al. (1998), we have

$$\|\nabla m(\mathbf{y})\| = \|\mathbf{y}\| \int_0^1 \kappa^{\frac{p}{2}+1} p(\kappa) e^{-\frac{Z}{2}\kappa} d\kappa$$

The score can be written as

$$\frac{\|\nabla m(\mathbf{y})\|}{p(\mathbf{y})} = \|\mathbf{y}\| \frac{m_{p+2}(\|\mathbf{y}\|)}{m_p(\|\mathbf{y}\|)} = \|\mathbf{y}\| E(\kappa | Z),$$

and the Laplacian term is $\Delta m(\mathbf{y}) = \int_0^1 (Z\kappa - p) \kappa^{\frac{p}{2}+1} p(\kappa) e^{-\frac{Z}{2}\kappa} d\kappa$. Combining these terms, we have,

$$\begin{aligned} \frac{\Delta m(\mathbf{y})}{p(\mathbf{y})} &= \frac{\int_0^1 (Z\kappa - p) \kappa^{\frac{p}{2}+1} p(\kappa) e^{-\frac{Z}{2}\kappa} d\kappa}{\int_0^1 \kappa^{\frac{p}{2}} p(\kappa) e^{-\frac{Z}{2}\kappa} d\kappa} \\ &= Z \frac{m_{p+4}(Z)}{m_p(Z)} - p \frac{m_{p+2}(Z)}{m_p(Z)}. \end{aligned}$$

The risk term $\Delta\sqrt{p(\mathbf{y})}/\sqrt{p(\mathbf{y})}$ is then computed using the identity

$$\frac{\nabla^2\sqrt{p(\mathbf{y})}}{\sqrt{p(\mathbf{y})}} = \frac{1}{2} \left[\frac{\Delta p(\mathbf{y})}{p(\mathbf{y})} - \frac{1}{2} \left\{ \frac{\|\nabla p(\mathbf{y})\|}{p(\mathbf{y})} \right\}^2 \right],$$

which reduces to

$$\frac{1}{2} \left\{ Z \frac{m_{p+4}(Z)}{m_p(Z)} - pg(Z) - \frac{Z}{2} g(Z)^2 \right\}$$

for $g(Z) = E(\kappa | Z)$.

Secondly, note that

$$Z\{m_{p+2}(Z) - m_{p+4}(Z)\} = 2 \int_0^1 \kappa^{\frac{p}{2}+1} (1-\kappa) p(\kappa) d\left(-e^{-\frac{Z}{2}\kappa}\right). \quad (16)$$

Therefore,

$$Z \left\{ \frac{m_{p+2}(Z)}{m_p(Z)} - \frac{m_{p+4}(Z)}{m_p(Z)} \right\} = \int_0^1 \left\{ (p+2)(1-\kappa) - 2\kappa + 2\kappa(1-\kappa) \frac{p'(\kappa)}{p(\kappa)} \right\} \frac{\kappa^{\frac{p}{2}} e^{-\frac{Z}{2}\kappa} p(\kappa)}{m_p(Z)} d\kappa$$

Then under the assumption that $\lim_{\kappa \rightarrow 0,1} \kappa(1-\kappa)p(\kappa) = 0$, integration by parts gives (8). Hence

$$E(\|\hat{\beta} - \beta\|^2) = p + 2E_{Z|\theta} \left[(Z+4)g(Z) - (p+2) - \frac{Z}{2}g(Z)^2 - E_{\kappa|Z} \left\{ 2\kappa(1-\kappa) \frac{p'(\kappa)}{p(\kappa)} \right\} \right].$$

□

References

- M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables*, volume 55 of *Applied Mathematics Series*. National Bureau of Standards, Washington, DC, 1964. Reprinted in paperback by Dover (1974); on-line at [http://www.math.sfu.ca/\\$sim\\$cbm/aands/](http://www.math.sfu.ca/simcbm/aands/).
- J. O. Berger. A robust generalized Bayes estimator and confidence region for a multivariate normal mean. *The Annals of Statistics*, 8(4):716–761, 1980.
- C. M. Carvalho, N. G. Polson, and J. G. Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–80, 2010.
- D. Fourdrinier, W. Strawderman, and M. T. Wells. On the construction of Bayes minimax estimators. *The Annals of Statistics*, 26(2):660–71, 1998.
- A. Gelman. Prior distributions for variance parameters in hierarchical models. *Bayesian Anal.*, 1(3):515–33, 2006.
- E. I. George, F. Liang, and X. Xu. Improved minimax predictive densities under Kullback-Leibler loss. *The Annals of Statistics*, 34(1):78–91, 2006.

- M. B. Gordy. A generalization of generalized beta distributions. Finance and Economics Discussion Series 1998-18, Board of Governors of the Federal Reserve System (U.S.), 1998.
- I. Gradshteyn and I. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, 1965.
- J. Griffin and P. Brown. Alternative prior distributions for variable selection with very many more variables than observations. Technical report, University of Warwick, 2005.
- H. Jeffreys. *Theory of Probability*. Oxford University Press, 3rd edition, 1961.
- Y. Maruyama. Improving on the James–Stein estimator. *Statistics and Decisions*, 14:137–40, 1999.
- Y. Maruyama and E. I. George. gbf: A fully Bayes factor with a generalized g-prior. Technical report, University of Tokyo, arXiv:0801.4410v2, 2010.
- C. Morris and R. Tang. Estimating random effects via adjustment for density maximization. *Statistical Science*, 26(2):271–87, 2011.
- N. G. Polson and J. G. Scott. Shrink globally, act locally: sparse Bayesian regularization and prediction. In *Proceedings of the 9th Valencia World Meeting on Bayesian Statistics*. Oxford Univeristy Press, 2011a.
- N. G. Polson and J. G. Scott. Local shrinkage rules, Lévy processes, and regularized regression. *Journal of the Royal Statistical Society (Series B)*, (to appear), 2011b.
- J. G. Scott and J. O. Berger. An exploration of aspects of Bayesian multiple testing. *Journal of Statistical Planning and Inference*, 136(7):2144–2162, 2006.
- C. Stein. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 9: 1135–51, 1981.
- W. Strawderman. Proper Bayes minimax estimators of the multivariate normal mean. *The Annals of Statistics*, 42:385–8, 1971.
- G. C. Tiao and W. Tan. Bayesian analysis of random-effect models in the analysis of variance. i. Posterior distribution of variance components. *Biometrika*, 51:37–53, 1965.
- R. Yang and J. O. Berger. A catalog of noninformative priors. Technical Report 42, Duke University Department of Statistical Science, 1997.